

End the Rebuild Cycle: The Agentic Knowledge Layer

Every new AI experience launches on infrastructure you've already built.

DATA SHEET

KEY AGENTIC KNOWLEDGE LAYER COMPONENTS

- ✓ Multi-layer indexing across vectors, full text and knowledge graphs
- ✓ 30+ configurable retrieval strategies, tunable without re-architecting downstream applications
- ✓ Sync agents that continuously ingest and refresh enterprise content
- ✓ REMi (RAG Evaluation Metrics Intelligence) for automated answer quality monitoring at production scale
- ✓ LLM-agnostic, SOC 2 Type II, ISO 27001 and GDPR-compliant from day one

Why Enterprise AI Investments Don't Compound

Building a retrieval-augmented generation (RAG) pipeline the first time is a solvable engineering problem. However, operating it in production is not. Ingestion failures, re-embedding jobs, retrieval tuning cycles and performance monitoring become standing responsibilities that expand in scope as usage grows. The engineering team that built the pipeline now owns it permanently, which means every hour spent on infrastructure maintenance is an hour not spent shipping AI features.

The second use case is where the economics break. A new AI assistant requires its own ingestion workflow, its own indexing configuration and its own retrieval logic. And if the underlying data overlaps with the first project, that data gets reprocessed and re-embedded from scratch. Vector stores duplicate, compute costs spike and governance controls get rewritten because there is no shared access policy layer to inherit. The operational surface area doubles, and so does the maintenance burden.

Hard-coded large language model (LLM) dependencies make it worse. When a model provider changes pricing, deprecates a version or a better-performing alternative enters the market, the teams that tied retrieval logic directly to application code pay the full cost of rework. Retrieval strategies, ranking logic and evaluation criteria are locked to infrastructure decisions made months earlier, under different constraints.

The Agentic Knowledge Layer in Practice

Stop rebuilding retrieval infrastructure for every use case. Start shipping AI experiences on a foundation that scales.

Build Once, Power Everything

One indexed knowledge layer eliminates duplicate ingestion, re-embedding and retrieval logic across projects. Then you can deploy a new AI experience without standing up another isolated stack or reprocessing data your team has already indexed.

Retrieval Quality You Can Measure, Not Assume

REMi (RAG Evaluation Metrics Intelligence) scores every answer for groundedness, context relevance and answer relevance at production scale, replacing manual spot checking with continuous automated evaluation. It allows you to catch retrieval failures before users do.

The RAG Stack That Pays for Itself as You Scale

The Agentic Knowledge Layer ingests once and reuses that foundation across every use case that follows, eliminating redundant compute costs across isolated pipelines. LLM-agnostic architecture lets teams route workloads to the most cost-efficient model available without rebuilding retrieval logic underneath.

Why the Progress® Agentic RAG Solution

Commodity RAG tools give you a retrieval primitive. Hyperscaler bundles lock retrieval to a model stack you don't control. The Agentic Knowledge Layer is purpose-built enterprise retrieval infrastructure that neither category can credibly offer.

- **Multi-layer indexing** across vectors, full text and knowledge graphs vs. single-strategy vector retrieval that misses context that commodity tools can't surface
- **REMi** continuous quality monitoring vs. no built-in evaluation framework, leaving retrieval quality to manual testing that doesn't scale
- **Governance, Role-Based Access Control (RBAC) and SOC 2 Type II, ISO 27001 and GDPR** compliance all built into the retrieval layer vs. access controls rebuilt from scratch for every new agent deployment
- **LLM-agnostic architecture** with SaaS, hybrid and on-premises deployment options vs. proprietary model stacks that make switching costs prohibitively high as the foundation model market reshuffles
- **30+ configurable retrieval strategies**, tunable without rearchitecting downstream applications vs. fixed retrieval logic hardcoded into application dependencies



Try the Progress Agentic RAG solution with a free 14-day trial today

About Progress Software

Progress Software (Nasdaq: PRGS) empowers organizations to achieve transformational success in the face of disruptive change. Our software enables our customers to develop, deploy and manage responsible AI-powered applications and personalized digital experiences with agility and ease. Businesses of all sizes get a trusted provider in Progress, with the products, expertise and vision they need to turn AI disruption into a competitive advantage. Millions of developers and technologists at hundreds of thousands of organizations depend on Progress every day. Learn more at www.progress.com

© 2026 Progress Software Corporation and/or its subsidiaries or affiliates.
All rights reserved. Rev 2026/07 | 1215969068920728

 /progresssw
 /progresssw
 /progresssw
 /progress-software
 /progress_sw_